



US009043594B2

(12) **United States Patent**
Sutton, II et al.

(10) **Patent No.:** **US 9,043,594 B2**
(45) **Date of Patent:** **May 26, 2015**

(54) **SYSTEM AND METHOD FOR EXECUTION OF A SECURED ENVIRONMENT INITIALIZATION INSTRUCTION**

(52) **U.S. Cl.**
CPC **G06F 9/4403** (2013.01); **G06F 9/44505** (2013.01); **G06F 12/1458** (2013.01); **G06F 21/57** (2013.01); **G01N 23/223** (2013.01); **G01N 33/502** (2013.01); **G01N 33/6872** (2013.01); **G01N 2223/076** (2013.01)

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)

(72) Inventors: **James A. Sutton, II**, Portland, OR (US); **David W. Grawrock**, Aloha, OR (US)

(58) **Field of Classification Search**
USPC 713/164
See application file for complete search history.

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(56) **References Cited**

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 14 days.

U.S. PATENT DOCUMENTS

(21) Appl. No.: **13/835,997**

5,123,101 A 6/1992 Sindhu
5,319,760 A 6/1994 Mason et al.
5,892,944 A 4/1999 Fukumoto et al.
6,035,380 A 3/2000 Shelton et al.

(22) Filed: **Mar. 15, 2013**

(Continued)

(65) **Prior Publication Data**

US 2013/0205127 A1 Aug. 8, 2013

OTHER PUBLICATIONS

Office Action received for U.S. Appl. No. 12/413,655, mailed on Apr. 13, 2010, 15 pages of Office Action.

(Continued)

Related U.S. Application Data

(63) Continuation of application No. 13/444,450, filed on Apr. 11, 2012, which is a continuation of application No. 12/455,844, filed on Jun. 8, 2009, now Pat. No. 8,185,734, which is a continuation of application No. 11/096,618, filed on Mar. 31, 2005, now Pat. No. 7,546,457, which is a continuation of application No. 10/112,169, filed on Mar. 29, 2002, now Pat. No. 7,069,442.

Primary Examiner — Matthew Smithers

(74) *Attorney, Agent, or Firm* — Thomas R. Lane


(51) **Int. Cl.**

H04L 29/00 (2006.01)
G06F 9/44 (2006.01)
G06F 9/445 (2006.01)
G06F 21/57 (2013.01)
G01N 23/223 (2006.01)
G01N 33/50 (2006.01)
G01N 33/68 (2006.01)
G06F 12/14 (2006.01)

(57) **ABSTRACT**

A method and apparatus for initiating secure operations in a microprocessor system is described. In one embodiment, one initiating logical processor initiates the process by halting the execution of the other logical processors, and then loading initialization and secure virtual machine monitor software into memory. The initiating processor then loads the initialization software into secure memory for authentication and execution. The initialization software then authenticates and registers the secure virtual machine monitor software prior to secure system operations.

3 Claims, 8 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

6,047,354 A 4/2000 Yoshioka et al.
6,075,938 A 6/2000 Bugnion et al.
6,327,646 B1 12/2001 Sharma et al.
6,356,989 B1 3/2002 Hays et al.
6,418,521 B1 7/2002 Mathews et al.
6,510,508 B1 1/2003 Zuraski et al.
6,604,187 B1 8/2003 McGrath et al.
6,779,085 B2 8/2004 Chauvel
7,069,442 B2 * 6/2006 Sutton et al. 713/179
7,409,487 B1 8/2008 Chen et al.
7,546,457 B2 * 6/2009 Sutton et al. 713/164
7,552,254 B1 6/2009 George et al.
7,552,255 B1 6/2009 George et al.
8,185,734 B2 * 5/2012 Sutton et al. 713/164
2002/0082824 A1 6/2002 Neiger et al.
2003/0217250 A1 11/2003 Bennett et al.

2004/0117593 A1 6/2004 Uhlig et al.
2004/0162930 A1 8/2004 Forin et al.
2007/0006230 A1 1/2007 Neiger et al.

OTHER PUBLICATIONS

Response to Office Action filed for U.S. Appl. No. 12/413,655, filed Jul. 9, 2010, 8 pages of Response to Office Action.
Office Action received for U.S. Appl. No. 12/483,519, mailed on Jul. 20, 2010, 17 pages of Office Action.
Office Action received for U.S. Appl. No. 12/413,655, mailed on Aug. 13, 2010, 12 Pages of Office Action.
Response to Office Action, for U.S. Appl. No. 12/413,655, filed Oct. 11, 2010, 6 Pages of Office Action.
Office Action received for Japanese Patent Divisional Application No. 2010-179219, mailed on Jan. 29, 2013, 3 pages of English Translation and 3 pages of Japanese Office Action.
Office Action received for Japanese Patent Divisional Application No. 2010-179219, mailed on May 28, 2013, 2 pages of English Translation and 3 pages of Japanese Office Action.

* cited by examiner




FIG. 1




FIG. 2

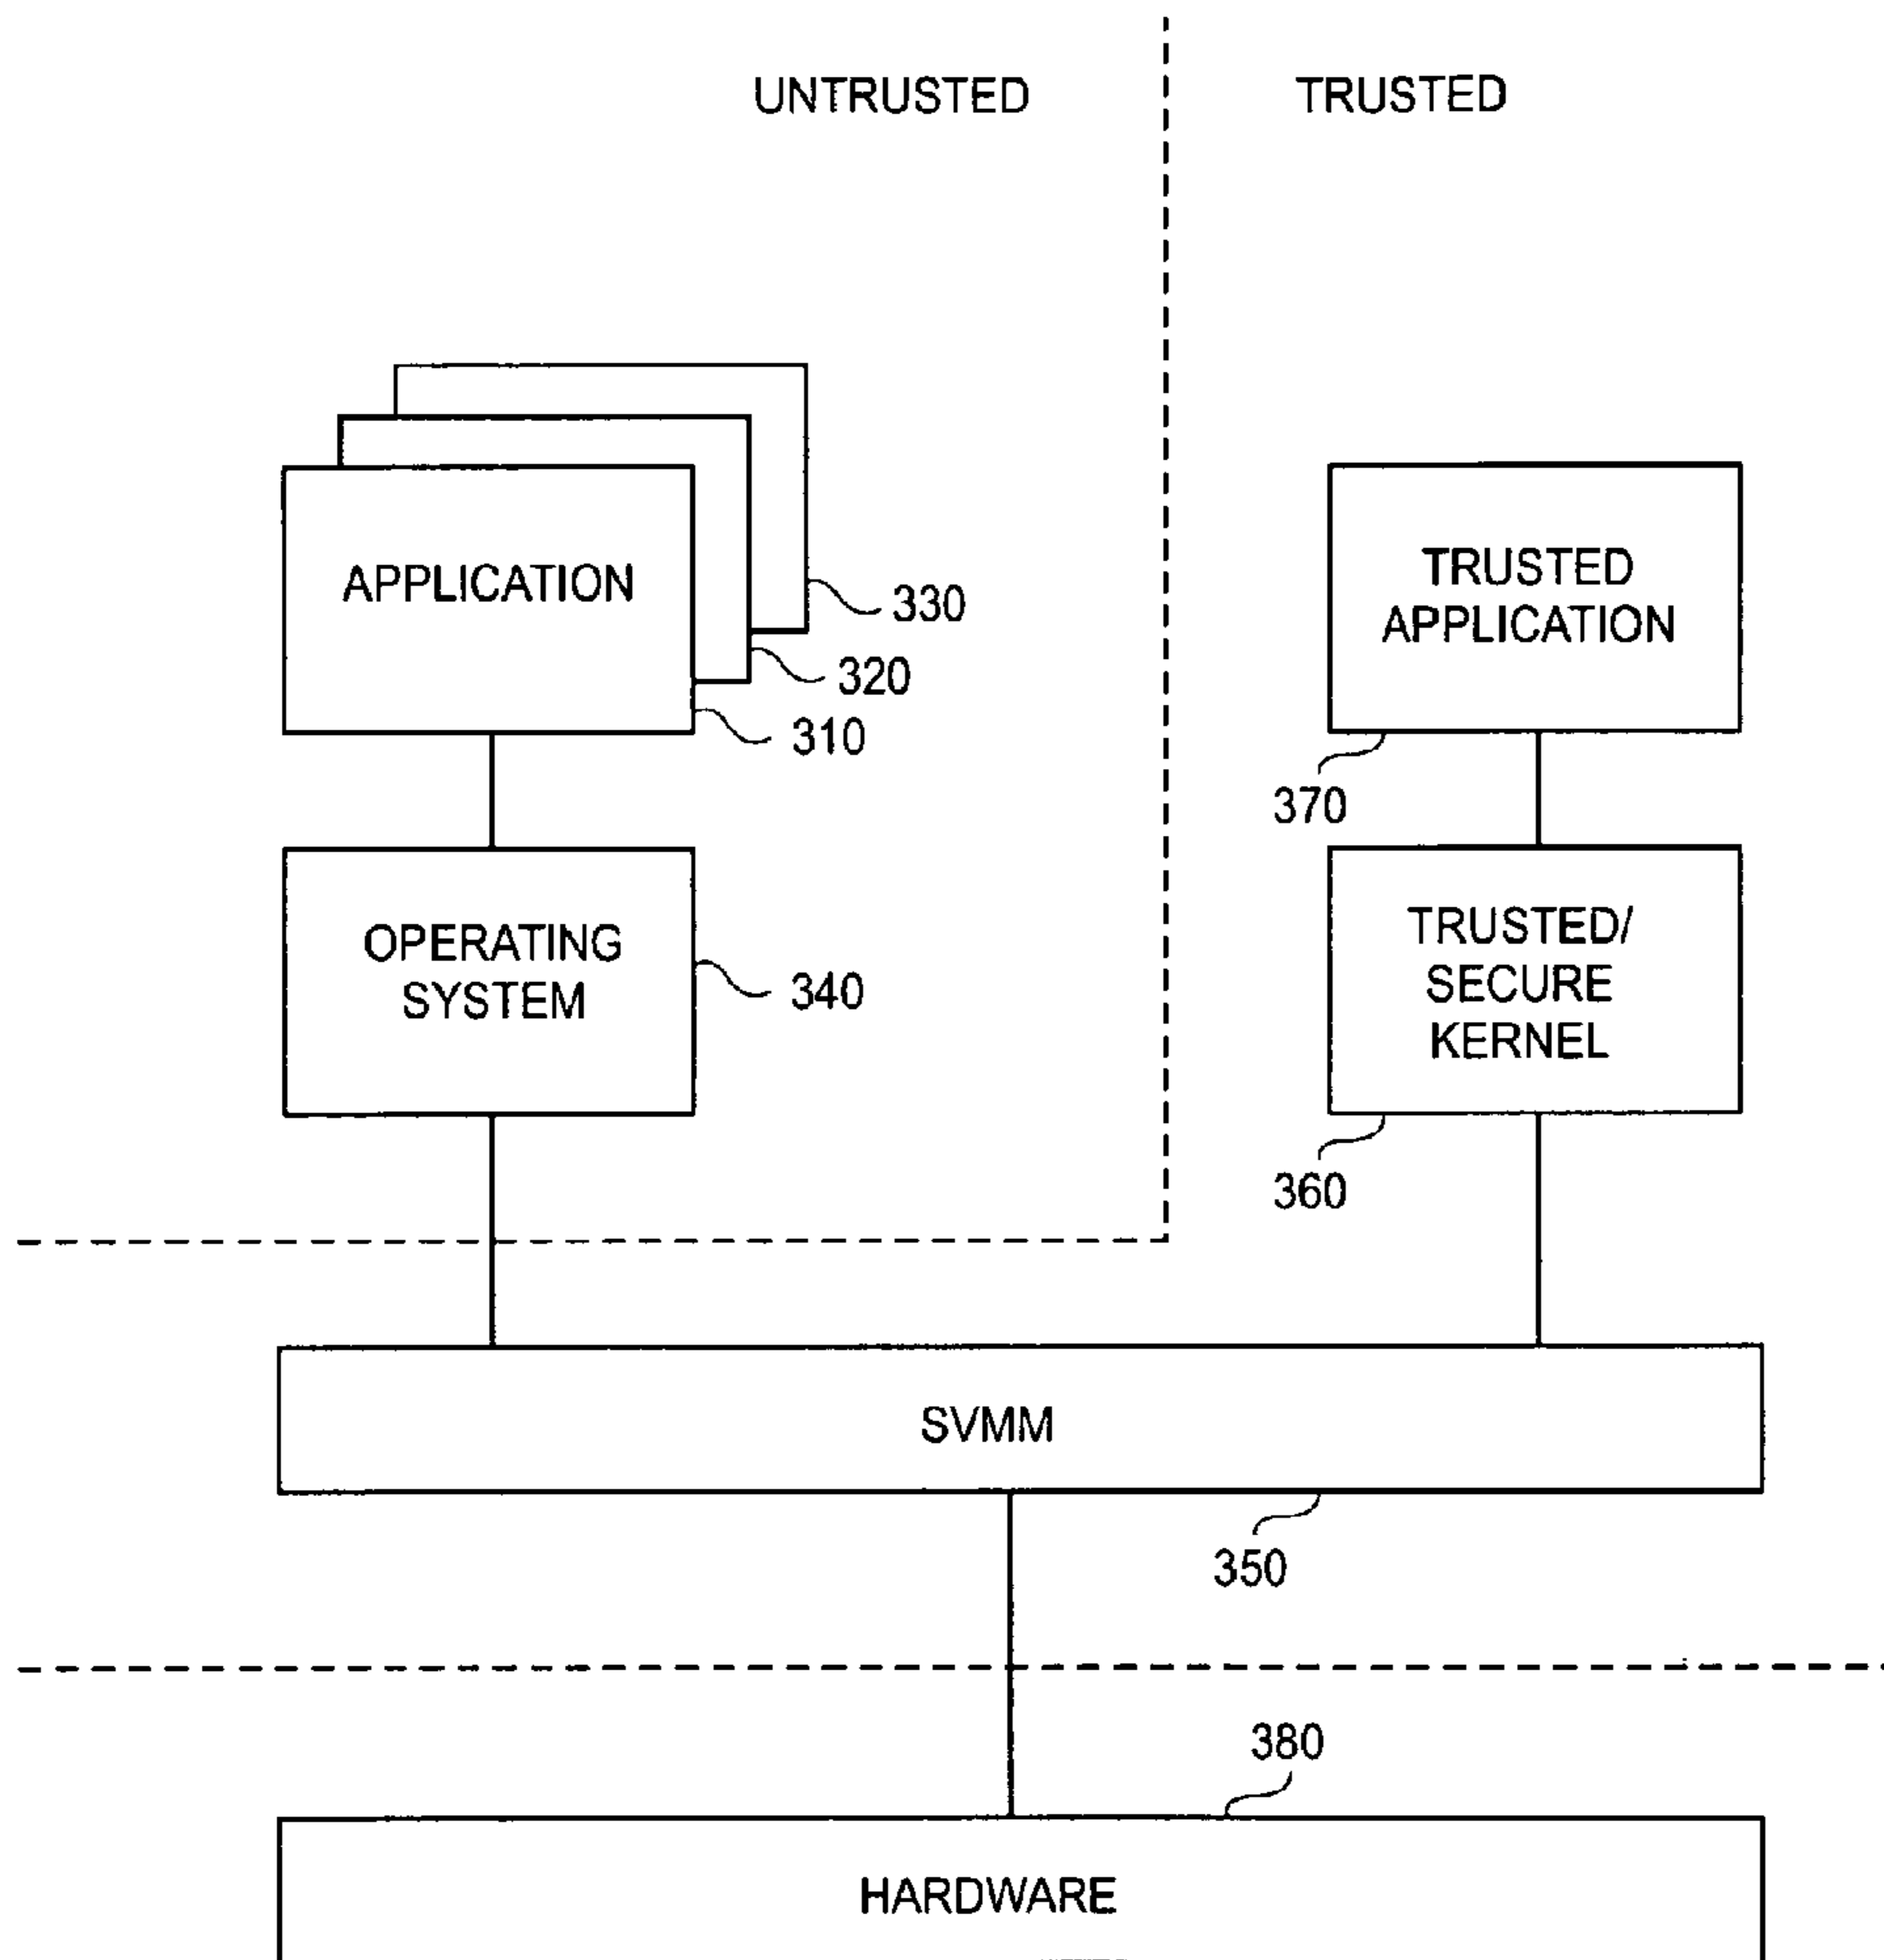


FIG. 3




FIG. 4A




FIG. 4B

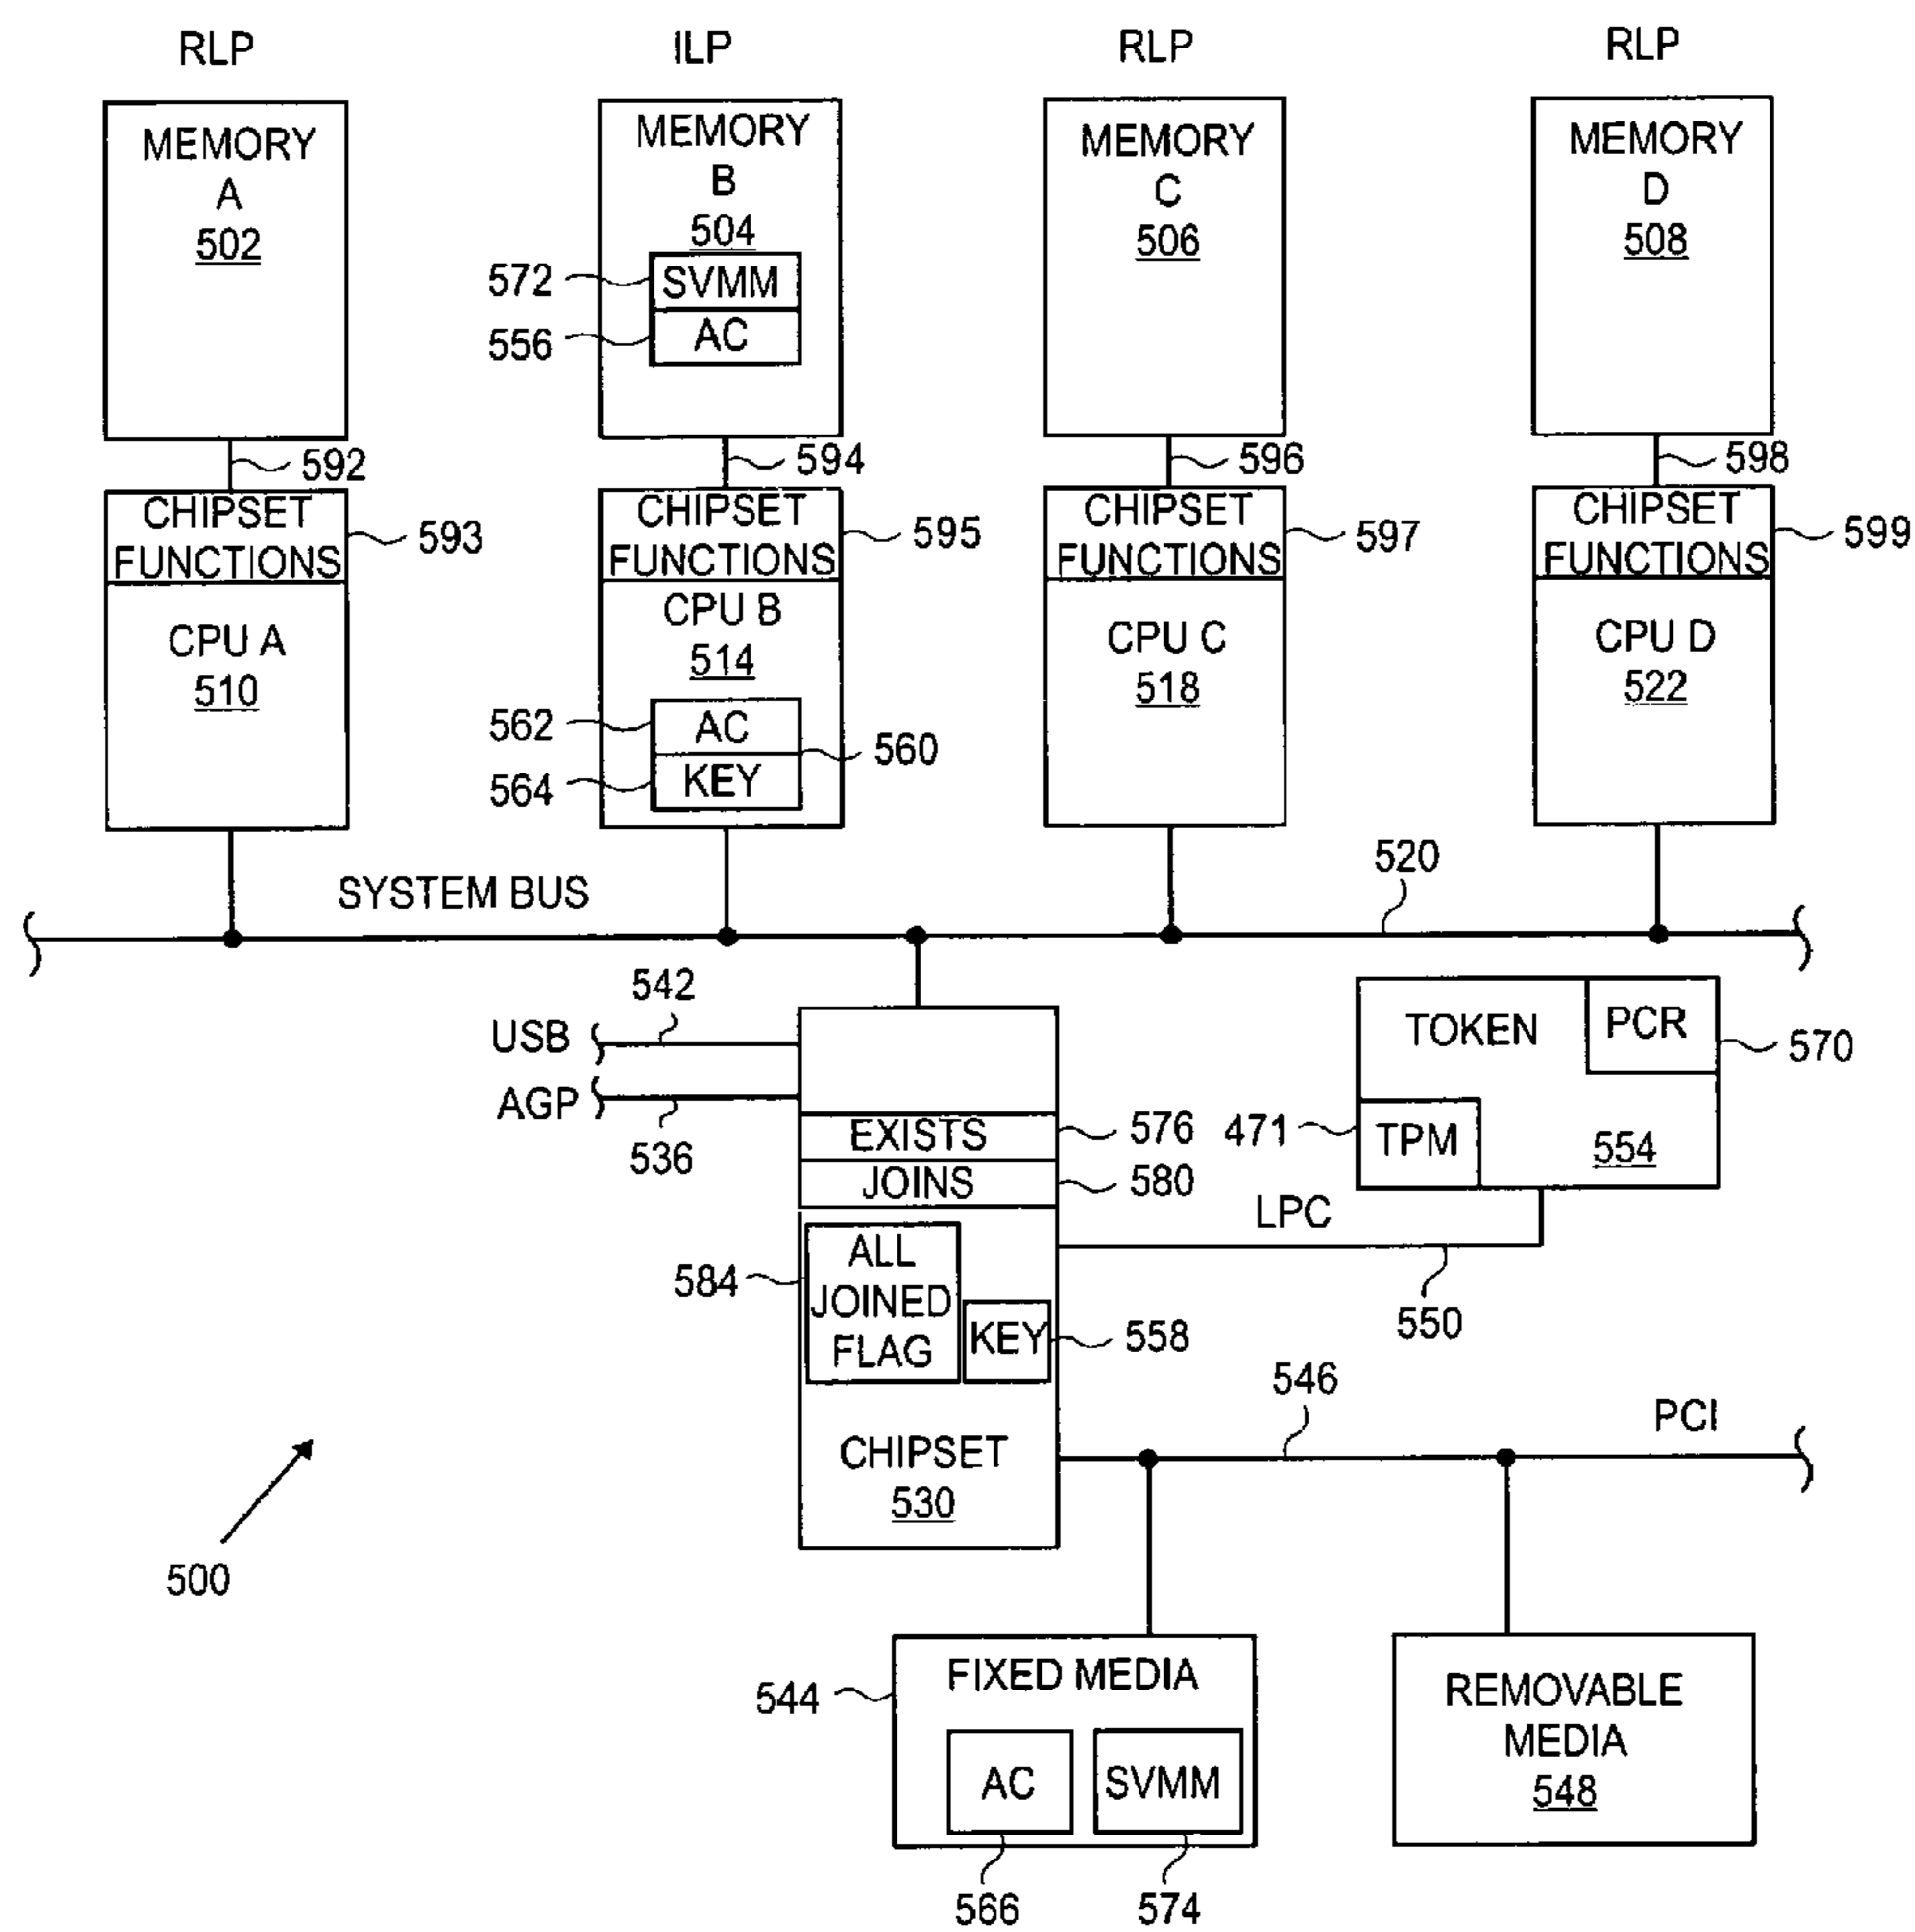


FIG. 5




FIG. 6

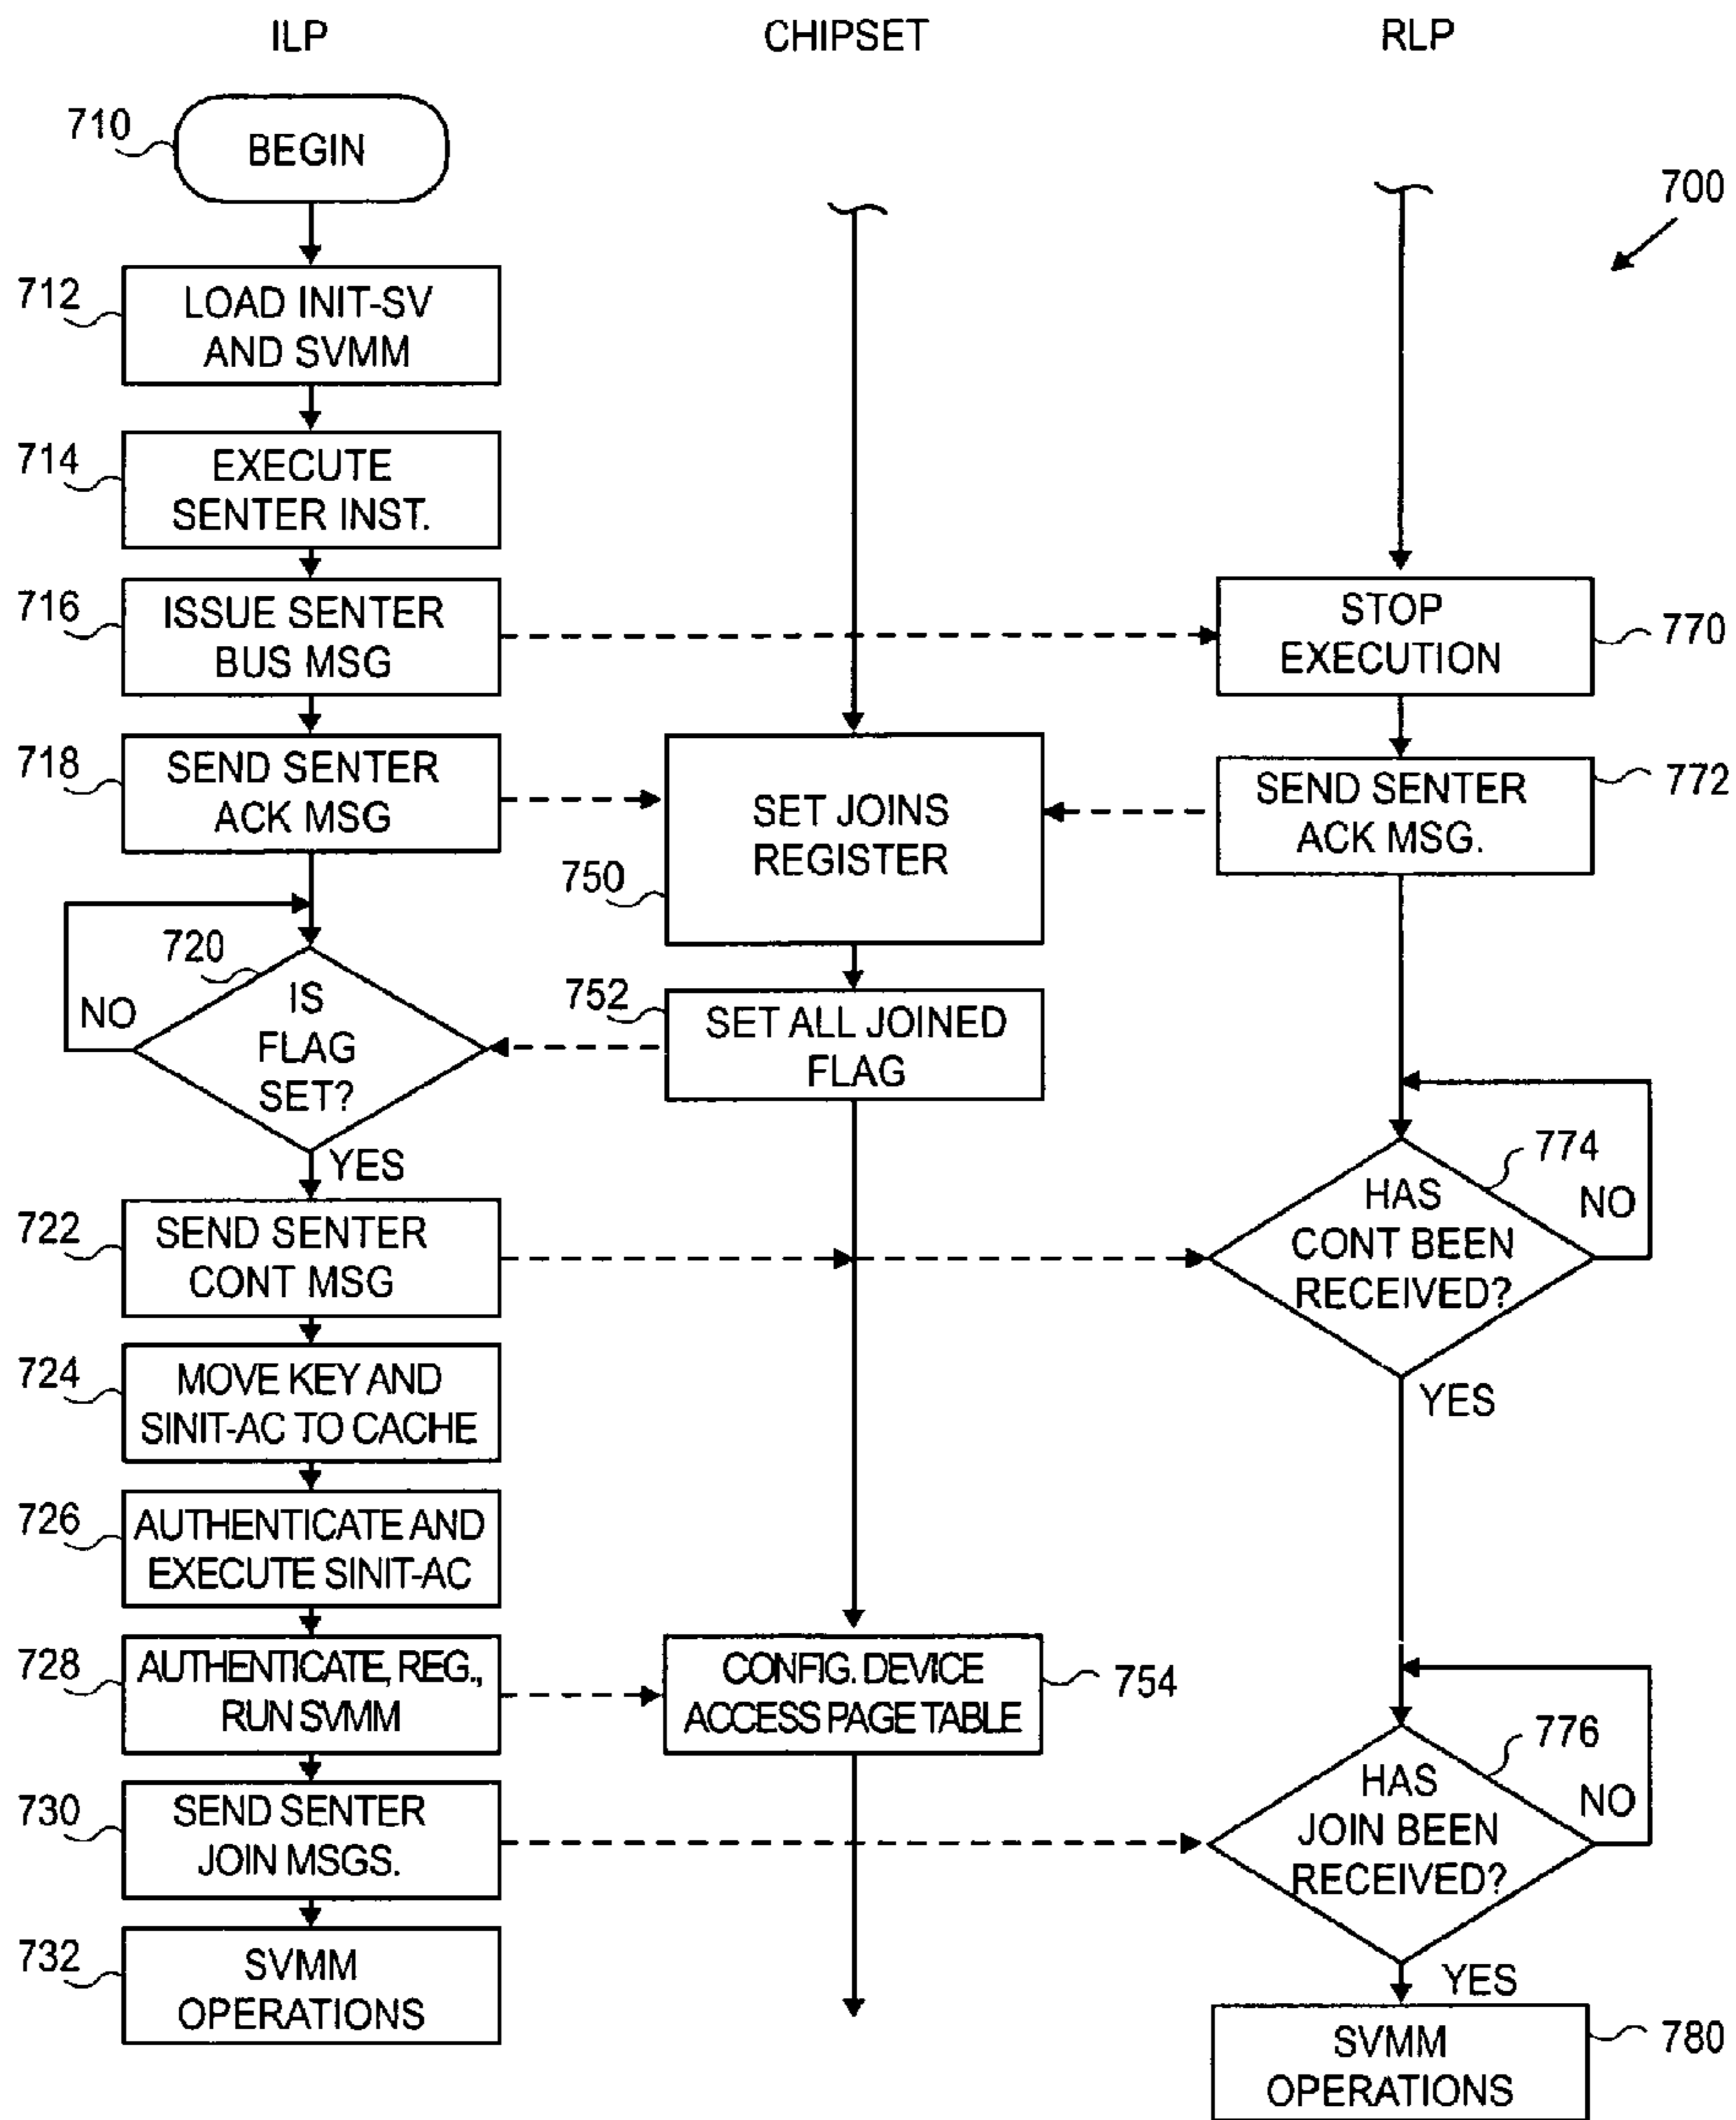


FIG. 7

1

SYSTEM AND METHOD FOR EXECUTION OF A SECURED ENVIRONMENT INITIALIZATION INSTRUCTION

RELATED APPLICATIONS

This application is a Continuation of U.S. patent application Ser. No. 13/444,450, entitled "SYSTEM AND METHOD FOR EXECUTION OF A SECURED ENVIRONMENT INITIALIZATION INSTRUCTION" filed on Apr. 11, 2012, which is a continuation of U.S. patent application Ser. No. 12/455,844, filed on Jun. 8, 2009, entitled "SYSTEM AND METHOD FOR EXECUTION OF A SECURED ENVIRONMENT INITIALIZATION INSTRUCTION" which issued on Jun. 9, 2009, as U.S. Pat. No. 8,185,734, which is a continuation of U.S. patent application Ser. No. 11/096,618, filed on Mar. 31, 2005, entitled "SYSTEM AND METHOD FOR EXECUTION OF A SECURED ENVIRONMENT INITIALIZATION INSTRUCTION" which issued on Jun. 9, 2009, as U.S. Pat. No. 7,546,457, is a continuation of U.S. patent application Ser. No. 10/112,169, filed on Mar. 29, 2002, entitled "SYSTEM AND METHOD FOR EXECUTION OF A SECURED ENVIRONMENT INITIALIZATION INSTRUCTION" which issued on Jun. 27, 2006, as U.S. Pat. No. 7,069,442.

FIELD

The present invention relates generally to microprocessor systems, and more specifically to microprocessor systems that may operate in a trusted or secured environment.

BACKGROUND

The increasing number of financial and personal transactions being performed on local or remote microcomputers has given impetus for the establishment of "trusted" or "secured" microprocessor environments. The problem these environments try to solve is that of loss of privacy, or data being corrupted or abused. Users do not want their private data made public. They also do not want their data altered or used in inappropriate transactions. Examples of these include unintentional release of medical records or electronic theft of funds from an on-line bank or other depository. Similarly, content providers seek to protect digital content (for example, music, other audio, video, or other types of data in general) from being copied without authorization.

Existing trusted systems may utilize a complete closed set of trusted software. This method is relatively simple to implement, but has the disadvantage of not allowing the simultaneous use of common, commercially available operating system and application software. This disadvantage limits the acceptance of such a trusted system.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is a diagram of an exemplary software environment executing in a microprocessor system.

FIG. 2 is a diagram of certain exemplary trusted or secured software modules and exemplary system environment, according to one embodiment of the present invention.

2

FIG. 3 is a diagram of an exemplary trusted or secured software environment, according to one embodiment of the present invention.

FIG. 4A is a schematic diagram of an exemplary microprocessor system adapted to support the secured software environment of FIG. 3, according to one embodiment of the present invention.

FIG. 4B is a schematic diagram of an exemplary microprocessor system adapted to support the secured software environment of FIG. 3, according to an alternate embodiment of the present invention.

FIG. 5 is a schematic diagram of an exemplary microprocessor system adapted to support the secured software environment of FIG. 3, according to an alternate embodiment of the present invention.

FIG. 6 is a time line drawing of the execution of software components, according to one embodiment of the present invention.

FIG. 7 is a flowchart of software and other process blocks, according to one embodiment of the present invention.

DETAILED DESCRIPTION

The following description describes techniques for initiating a trusted or secured environment in a microprocessor system. In the following description, numerous specific details such as logic implementations, software module allocation, encryption techniques, bus signaling techniques, and details of operation are set forth in order to provide a more thorough understanding of the present invention. It will be appreciated, however, by one skilled in the art that the invention may be practiced without such specific details. In other instances, control structures, gate level circuits and full software instruction sequences have not been shown in detail in order not to obscure the invention. Those of ordinary skill in the art, with the included descriptions, will be able to implement appropriate functionality without undue experimentation. The invention is disclosed in the form of a microprocessor system. However, the invention may be practiced in other forms of processor such as a digital signal processor, a mini-computer, or a mainframe computer.

Referring now to FIG. 1, a diagram of an exemplary software environment executing in a microprocessor system is shown. The software shown in FIG. 1 is not trusted (untrusted). When operating in a high privilege level, the size and constant updating of the operating system 150 make it very difficult to perform any trust analysis in a timely manner. Much of the operating system sits within privilege ring zero (0), the highest level of privilege. The applications 152, 154, and 156 have much reduced privilege and typically reside within privilege ring three (3). The existence of the differing privilege rings and the separation of the operating system 150 and applications 152, 154 and 156 into these differing privileged rings would appear to allow operating of the software of FIG. 1 in a trusted mode, based on making a decision to trust the facilities provided by the operating system 150. However, in practice making such a trust decision is often impractical. Factors that contribute to this problem include the size (number of lines of code) of the operating system 150, the fact that the operating system 150 may be the recipient of numerous updates (new code modules and patches) and the fact that the operating system 150 may also contain code modules such as device drivers supplied by parties other than the operating system developer. Operating system 150 may be a common one such as Microsoft® Windows®, Linux, or Solaris®, or may be any other appropriate known or otherwise available

operating system. The particular types or names of applications or operating systems run or running are not critical.

Referring now to FIG. 2, a diagram of certain exemplary trusted or secured software modules and exemplary system environment 200 is shown, according to one embodiment of the present invention. In the FIG. 2 embodiment, processor 202, processor 212, processor 222, and optional other processors (not shown) are shown as separate hardware entities. In other embodiments, the number of processors may differ, as may the boundary of various components and functional units. In some embodiments the processors may be replaced by separate hardware execution threads or "logical processors" running on one or more physical processors.

Processors 202, 212, 222 may contain certain special circuits or logic elements to support secure or trusted operations. For example, processor 202 may contain secure enter (SENDER) logic 204 to support the execution of special SENTER instructions that may initiate trusted operations. Processor 202 may also contain bus message logic 206 to support special bus messages on system bus 230 in support of special SENTER operations. In alternate embodiments, memory control functions of chipset 240 may be allocated to circuits within the processors, and for multiple processors may be included on a single die. In these embodiments, special bus messages may also be sent on busses internal to the processors. The use of special bus messages may increase the security or trustability of the system for several reasons. Circuit elements such as processors 202, 212, and 222 or chipset 240 may only issue or respond to such messages if they contain the appropriate logic elements of embodiments of the present disclosure. Therefore successful exchange of the special bus messages may help ensure proper system configuration. Special bus messages may also permit activities that should normally be prohibited, such as resetting a platform configuration register 278. The ability of potentially hostile untrusted code to spy on certain bus transactions may be curtailed by allowing special bus messages to be issued only in response to special security instructions.

Additionally, processor 202 may contain secure memory 208 to support secure initialization operations. In one embodiment secure memory 208 may be an internal cache of processor 202, perhaps operating in a special mode. In alternate embodiments secure memory 208 may be special memory. Other processors such as processor 212 and processor 222 may also include SENTER logic 214, 224, bus message logic 216, 226, and secure memory 218, 228.

A "chipset" may be defined as a group of circuits and logic that support memory and input/output (I/O) operations for a connected processor or processors. Individual elements of a chipset may be grouped together on a single chip, a pair of chips, or dispersed among multiple chips, including processors. In the FIG. 2 embodiment, chipset 240 may include circuitry and logic to support memory and I/O operations to support processors 202, 212, and 222. In one embodiment, chipset 240 may interface with a number of memory pages 250 through 262 and a device-access page table 248 containing control information indicating whether non-processor devices may access the memory pages 250 through 262. Chipset 240 may include device-access logic 247 that may permit or deny direct memory access (DMA) from I/O devices to selected portions of the memory pages 250 through 262. In some embodiment the device access logic 247 may contain all relevant information required to permit or deny such accesses. In other embodiments, the device access logic 247 may access such information held in the device access page table 248. The actual number of memory pages is not important and will change depending upon system require-

ments. In other embodiments the memory access functions may be external to chipset 240. The functions of chipset 240 may further be allocated among one or more physical devices in alternate embodiments.

Chipset 240 may additionally include its own bus message logic 242 to support special bus messages on system bus 230 in support of special SENTER operations. Some of these special bus messages may include transferring the contents of a key register 244 to a processor 202, 212, or 222, or permitting a special ALL_JOINED flag 274 to be examined by a processor 202, 212, or 222. Additional features of the bus message logic 242 may be to register bus activity by processors in an "EXISTS" register 272 and store certain special bus message activity by processors in a "JOINS" register 272. Equality of contents of EXISTS register 272 and JOINS register 272 may be used to set the special ALL_JOINED flag 274 to indicate all processors in the system are participating in the secure enter process.

Chipset 240 may support standard I/O operations on I/O busses such as peripheral component interconnect (PCI), accelerated graphics port (AGP), universal serial bus (USB), low pin count (LPC) bus, or any other kind of I/O bus (not shown). An interface 290 may be used to connect chipset 240 with token 276, containing one or more platform configuration registers (PCR) 278, 279. In one embodiment, interface 290 may be the LPC bus (Low Pin Count (LPC) Interface Specification, Intel Corporation, rev. 1.0, 29 Dec. 1997) modified with the addition of certain security enhancements. One example of such a security enhancement would be a locality confirming message, utilizing a previously-reserved message header and address information targeting a platform configuration register (PCR) 278 within token 276. In one embodiment, token 276 may contain special security features, and in one embodiment may include the trusted platform module (TPM) 281 disclosed in the Trusted Computing Platform Alliance (TCPA) Main Specification, version 1.1a, 1 Dec. 2001, issued by the TCPA (available at the time of filing of the present application at www.trustedpc.com).

Two software components identified in system environment 200 are a Secure Virtual Machine Monitor (SVMM) 282 module and a Secure Initialization Authenticated Code (SINIT-AC) 280 module. The SVMM 282 module may be stored on a system disk or other mass storage, and moved or copied to other locations as necessary. In one embodiment, prior to beginning the secure launch process SVMM 282 may be moved or copied to one or more memory pages 250 through 262. Following the secure enter process, a virtual machine environment may be created in which the SVMM 282 may operate as the most privileged code within the system, and may be used to permit or deny direct access to certain system resources by the operating system or applications within the created virtual machines.

Some of the actions required by the secure enter process may be beyond the scope of simple hardware implementations, and may instead advantageously use a software module whose execution can be implicitly trusted. In one embodiment, these actions may be performed by Secure Initialization (SINIT) code. Three exemplary actions are identified here, but these actions should not be taken to be limiting. One action may require that various controls representing critical portions of the system configuration be tested to ensure that the configuration supports the correct instantiation of the secure environment. In one embodiment, one required test may be that the memory controller configuration provided by chipset 240 does not permit two or more different system bus addresses to touch the same location within memory pages 250 through 262. A second action may be to configure the

device-access page table **248** and device-access logic **247** to protect those memory pages used by the memory-resident copy of SVMM **282** from interference by non-processor devices. A third action may be to calculate and register the SVMM **282** module's identity and transfer system control to it. Here "register" means placing a trust measurement of SVMM **282** into a register, for example into PCR **278** or into PCR **279**. When this last action is taken, the trustworthiness of the SVMM **282** may be inspected by a potential system user.

The SINIT code may be produced by the manufacturer of the processors or of the chipsets. For this reason, the SINIT code may be trusted to aid in the secure launch of chipset **240**. In order to distribute the SINIT code, in one embodiment a well-known cryptographic hash is made of the entire SINIT code, producing a value known as a "digest". One embodiment produces a 160-bit value for the digest. The digest may then be encrypted by a private key, held in one embodiment by the manufacturer of the processor, to form a digital signature. When the SINIT code is bundled with the corresponding digital signature, the combination may be referred to as SINIT authenticated code (SINIT-AC) **280**. Copies of the SINIT-AC **280** may be later validated as discussed below.

The SINIT-AC **280** may be stored on system disk or other mass storage or in a fixed media, and moved or copied to other locations as necessary. In one embodiment, prior to beginning the secure launch process SINIT-AC **280** may be moved or copied into memory pages **250-262** to form a memory-resident copy of SINIT-AC.

Any logical processor may initiate the secure launch process, and may then be referred to as the initiating logical processor (ILP). In the present example processor **202** becomes the ILP, although any of the processors on system bus **230** could become the ILP. Neither memory-resident copy of SINIT-AC **280** nor memory-resident copy of SVMM **282** may be considered trustworthy at this time since, among other reasons, the other processors or the DMA devices may overwrite memory pages **250-262**.

The ILP (processor **202**) then executes a special instruction. This special instruction may be referred to as a secured enter (SENDER) instruction, and may be supported by SENDER logic **204**. Execution of the SENDER instruction may cause the ILP (processor **202**) to issue special bus messages on system bus **230**, and then wait considerable time intervals for subsequent system actions. After execution of SENDER begins, one of these special bus messages, SENDER BUS MESSAGE, is broadcast on system bus **230**. Those logical processors other than the ILP, which may be referred to as responding logical processors (RLPs), respond to the SENDER BUS MESSAGE with an internal non-maskable event. In the present example, the RLPs include processor **212** and processor **222**. The RLPs must each terminate current operations, send a RLP acknowledge (ACK) special bus message on system bus **230**, and then enter a wait state. It should be noted that the ILP also sends its own ACK message over system bus **230**.

The chipset **240** may contain a pair of registers, "EXISTS" register **270** and "JOINS" register **272**. These registers may be used to verify that the ILP and all of the RLPs are responding properly to the SENDER BUS MESSAGE. In one embodiment, chipset **240** may keep track of all operational logical processors in the system by writing a "1" into the corresponding bit of the EXISTS register **270** on any system bus transaction made by that logical processor. In this embodiment, each transaction on system bus **230** must contain an identification field containing the logical processor identifier. In one embodiment, this consists of a physical processor identifier and an identifier for the hardware execu-

tion thread within each physical processor. For example, if a thread executing on processor **222** caused any bus transactions on system bus **230**, chipset **240** would see this logical processor identifier within the transaction and write a "1" into the corresponding location **286** within EXISTS register **270**. During the secure launch process, when that same thread on processor **222** sends its ACK message on system bus **230**, the chipset **240** would also see this and could write a "1" into the corresponding location **288** in the JOINS register **272**. (In the FIG. 2 example, each physical processor is shown with only a single thread executing for clarity. In alternate embodiments the physical processors may support multiple threads, and thereby multiple logical processors.) When the contents of the JOINS register **272** matches the contents of the EXISTS register **270**, then chipset **240** can set an ALL_JOINED flag **246** indicating that all processors have properly responded to the SENDER BUS MESSAGE.

In another embodiment, EXISTS register **270** and JOINS register **272** may continue to aid security subsequent to the setting of the ALL_JOINED flag **246**. During the time subsequent to the setting of the ALL_JOINED flag **246** until the end of trusted or secure operations, chipset **240** may continue to monitor and compare bus cycles against the JOINS register **272**. During this period, if chipset **240** ever sees a bus transaction from a logical processor that is not currently identified in JOINS register **272**, then chipset **240** may presume that this logical processor has somehow "appeared" late. This would imply that such a logical processor did not participate in the secure launch process, and therefore could represent an attacker (security threat). In such circumstances, chipset **240** may respond appropriately to keep this attacker out of the secured environment. In one embodiment, chipset **240** may force a system reset in such circumstances. In a second embodiment, similar detection of a "late" processor may be achieved by each logical processor asserting a special reserved signal on the system bus on every transaction following the assertion of the ACK bus message. In this embodiment, following the setting of the ALL_JOINED flag **246** if the chipset **240** observes a bus transaction initiated by a processor without the special signal asserted, then chipset **240** may again presume that this logical processor has somehow appeared "late", and may represent an attacker.

After issuing the SENDER BUS MESSAGE, the ILP (processor **202**) polls the ALL_JOINED flag **246** to see when and if all processors have properly responded with their ACKs. If the flag **246** is never set, several implementations are possible. A watchdog timer in the ILP or chipset or elsewhere may cause a system reset. Alternatively, the system may hang requiring operator reset. In either case the assertion of a secure environment is protected (in that the secure launch process does not complete unless all processors participate), although the system may not continue to function. In normal operations, after a short time the ALL_JOINED flag **246** is set, and the ILP may be assured that all other logical processors have entered a wait state.

When the ALL_JOINED flag **246** is set, the ILP (processor **202**) may move both a copy of SINIT-AC **280** and key **284** into secure memory **208** for the purpose of authenticating and subsequently executing the SINIT code included in SINIT-AC **280**. In one embodiment, this secure memory **208** may be an internal cache of the ILP (processor **202**), perhaps operating in a special mode. Key **284** represents the public key corresponding to the private key used to encrypt the digital signature included in the SINIT-AC **280** module, and is used to verify the digital signature and thereby authenticate the SINIT code. In one embodiment, key **284** may already be stored in the processor, perhaps as part of the SENDER logic

204. In another embodiment, key 284 may be stored in a read-only key register 244 of chipset 240, which is read by the ILP. In yet another embodiment, either the processor or the chipset's key register 244 may actually hold a cryptographic digest of key 284, where key 284 itself is included in the SINIT-AC 280 module. In this last embodiment, the ILP reads the digest from key register 244, calculates an equivalent cryptographic hash over the key 284 embedded in SINIT-AC 280, and compares the two digests to ensure the supplied key 284 is indeed trusted.

A copy of SINIT-AC and a copy of a public key may then exist within secure memory 208. The ILP may now validate the copy of SINIT-AC by decrypting the digital signature included in the copy of the SINIT-AC using the copy of a public key. This decryption produces an original copy of a cryptographic hash's digest. If a newly-calculated digest matches this original digest then the copy of SINIT-AC and its included SINIT code may be considered trustable.

The ILP may now issue another special bus message, SENTER CONTINUE MESSAGE, via system bus 230 signaling the waiting RLP's (processor 212, processor 222) and chipset 240 that secured operations are going to be initiated. The ILP may now register the unique identity of the SINIT-AC module by writing the SINIT-AC module's cryptographic digest value to a platform configuration register 272 in the security token 276, as outlined below. The ILP's execution of its SENTER instruction may now terminate by transferring execution control to the trusted copy of the SINIT code held within the ILP's secure memory 208. The trusted SINIT code may then perform its system test and configuration actions and may register the memory-resident copy of SVMM, in accordance with the definition of "register" above.

Registration of the memory-resident copy of SVMM may be performed in several manners. In one embodiment, the SENTER instruction running on the ILP writes the calculated digest of SINIT-AC into PCR 278 within the security token 276. Subsequently, the trusted SINIT code may write the calculated digest of the memory-resident SVMM to the same PCR 278 or another PCR 279 within the security token 276. If the SVMM digest is written to the same PCR 278, the security token 276 hashes the original contents (SINIT digest) with the new value (SVMM digest) and writes the result back into the PCR 278. In embodiments where the first (initializing) write to PCR 278 is limited to the SENTER instruction, the resulting digest may be used as a root of trust for the system.

Once the trusted SINIT code has completed its execution, and has registered the identity of the SVMM in a PCR, the SINIT code may transfer ILP execution control to the SVMM. In a typical embodiment, the first SVMM instructions executed by the ILP may represent a self-initialization routine for the SVMM. The ILP may in one embodiment issue individual RLP JOIN MESSAGE special bus messages to each RLP, causing each of the RLPs to join in operations under the supervision of the now-executing copy of SVMM. From this point onwards, the overall system is operating in trusted mode as outlined in the discussion of FIG. 3 below.

Referring now to FIG. 3, a diagram of an exemplary trusted or secured software environment is shown, according to one embodiment of the present invention. In the FIG. 3 embodiment, trusted and untrusted software may be loaded simultaneously and may execute simultaneously on a single computer system. A SVMM 350 selectively permits or prevents direct access to hardware resources 380 from one or more untrusted operating systems 340 and untrusted applications 310 through 330. In this context, "untrusted" does not necessarily mean that the operating system or applications are

deliberately misbehaving, but that the size and variety of interacting code makes it impractical to reliably assert that the software is behaving as desired, and that there are no viruses or other foreign code interfering with its execution. In a typical embodiment, the untrusted code might consist of the normal operating system and applications found on today's personal computers.

SVMM 350 also selectively permits or prevents direct access to hardware resources 380 from one or more trusted or secure kernels 360 and one or more trusted applications 370. Such a trusted or secure kernel 360 and trusted applications 370 may be limited in size and functionality to aid in the ability to perform trust analysis upon it. The trusted application 370 may be any software code, program, routine, or set of routines which is executable in a secure environment. Thus, the trusted application 370 may be a variety of applications, or code sequences, or may be a relatively small application such as a Java applet.

Instructions or operations normally performed by operating system 340 or kernel 360 that could alter system resource protections or privileges may be trapped by SVMM 350, and selectively permitted, partially permitted, or rejected. As an example, in a typical embodiment, instructions that change the processor's page table that would normally be performed by operating system 340 or kernel 360 would instead be trapped by SVMM 350, which would ensure that the request was not attempting to change page privileges outside the domain of its virtual machine.

Referring now to FIG. 4A, one embodiment of a microprocessor system 400 adapted to support the secured software environment of FIG. 3 is shown. CPU A 410, CPU B 414, CPU C 418, and CPU D 422 may be configured with additional microcode or logic circuitry to support the execution of special instructions. In one embodiment, this additional microcode or logic circuitry may be the SENTER logic 204 of FIG. 2. These special instructions may support the issuance of special bus messages on system bus 420 that may enable the proper synchronization of the processors while launching the secure environment. In one embodiment, the issuance of special bus messages may be supported by circuitry such as the bus message logic 206 of FIG. 2. Similarly chipset 430 may be similar to chipset 240 and may support the above-mentioned special cycles on system bus 420. The number of physical processors may vary upon the implementation of a particular embodiment. In one embodiment, the processors may be Intel® Pentium® class microprocessors. Chipset 430 may interface with mass storage devices such as fixed media 444 or removable media 448 via PCI bus 446, or, alternately, via USB 442, an integrated controller electronics (IDE) bus (not shown), a small computer systems interconnect (SCSI) bus (not shown), or any other I/O busses. The fixed media 444 or removable media 448 may be magnetic disks, magnetic tape, magnetic diskettes, magneto-optical drives, CD-ROM, DVD-ROM, Flash memory cards, or many other forms of mass storage.

In the FIG. 4A embodiment, the four processors CPU A 410, CPU B 414, CPU C 418, and CPU D 422 are shown as four separate hardware entities. In other embodiments, the number of processors may differ. Indeed, the physically discrete processors may be replaced by separate hardware execution threads running on one or more physical processors. In the latter case these threads possess many of the attributes of additional physical processors. In order to have a generic expression to discuss using any mixture of multiple physical processors and multiple threads upon processors, the expression "logical processor" may be used to describe either a physical processor or a thread operating in one or more

physical processors. Thus, one single-threaded processor may be considered a logical processor, and multi-threaded or multi-core processors may be considered multiple logical processors.

In one embodiment, chipset **430** interfaces with a modified LPC bus **450**. Modified LPC bus **450** may be used to connect chipset **430** with a security token **454**. Token **454** may in one embodiment include the TPM **471** envisioned by the Trusted Computing Platform Alliance (TCPA).

Referring now to FIG. **4B**, an alternate embodiment of a microprocessor system **490** adapted to support the secured software environment of FIG. **3** is shown. Differing from the FIG. **4A** embodiment, CPU A **410** and CPU B **414** may be connected to chipset **428** with system bus A **402** whereas CPU C **418** and CPU D **422** may be connected to chipset **428** with system bus B **404**. In other embodiments more than two system busses may be utilized. In another alternative embodiment, point-to-point busses may be used. Special instructions may support the issuance of special bus messages on system bus A **402** and system bus B **404** that may enable the proper synchronization of the processors while launching the secure environment. In one embodiment, the issuance of special bus messages may be supported by circuitry such as the bus message logic **206** of FIG. **2**.

In one embodiment, chipset **428** is responsible for maintaining consistency and coherency across system bus A **402** and system bus B **404**. If a bus message, standard or special, is sent across system bus A **402**, chipset **428** reflects that message (when appropriate) onto system bus B **404**, and vice-versa.

In an alternate embodiment, chipset **428** treats system bus A **402** and system bus B **404** as independent subsystems. Any special bus messages issued on system bus A **402** apply only to processors on that bus: similarly, special bus messages issued on system bus B **404** apply only to processors on that bus. Any protected memory that is established with respect to system bus A **402** is only accessible to processors connected to system bus A **402**, and the processors on system bus B **404** may be treated as untrusted devices. To gain access to any protected memory established for CPU A **410** and CPU B **414** on system bus A **402**, processors CPU C **418** and CPU D **422** on system bus B **404** must perform their own SENTER process, creating a registered environment equal to that created for the processors on system bus A **402**.

Referring now to FIG. **5**, a schematic diagram of an exemplary microprocessor system **500** adapted to support the secured software environment of FIG. **3** is shown, according to an alternate embodiment of the present invention. Differing from the FIG. **4A** embodiment, each processor (for example, CPU A **510**) may include certain chipset functions (for example, chipset functions **593**) that, for example, perform memory controller functions and device access logic functions. These chipset functions thereby allow the direct connection of memory (for example, memory A **502**) to the processor. Other chipset functions may remain in a separate chipset **530**. Special bus messages may be issued across system bus **520**.

Each processor may make indirect accesses to memory connected to other processors: however, these accesses may be considerably slower when compared to accesses to a processor's own memory. Prior to the start of the SENTER process, software may move copies of SINIT-AC **566** and SVMM **574** from fixed media **544** into local memory **504**, forming copy of SINIT-AC **556** and copy of SVMM **572**. In one embodiment, the memory **504** may be selected because it is directly accessed by the processor intended to be the ILP, in the FIG. **5** example this is CPU B **514**. Alternatively, the

SINIT-AC **566** and SVMM **574** copies may be placed in other memories attached to other (non-ILP) processors, so long as the ILP **514** has the ability to access those memories. CPU B ILP **514** begins the secure enter process by issuing the SENTER instruction, as already described in FIG. **2**, and with similar consequences and bus cycles issued. Chipset **530** may utilize EXISTS register **576**, JOINS register **580**, and ALL_JOINED flag **584** as described above in connection with FIG. **2** to determine whether all processors have properly responded to the SENTER BUS MESSAGE and signal this information to the ILP. The ILP (CPU B **514**) may again move the memory-resident copy of SINIT-AC **556** into secure memory **560**, along with a copy of a public key **564**. Upon verification and registration of SINIT-AC **556**, ILP may then continue to verification and registration of the memory-resident copy of SVMM **572**.

Referring now to FIG. **6**, a time line drawing of various operations is shown, according to one embodiment of the present invention. The timeline of FIG. **6** shows the overall schedule of the operations discussed in connection with the exemplary system discussed in connection with FIG. **2** above. When software decides that secure or trusted operations are desired, at time **610** any software locates and makes a copy of SINIT-AC **280** and SVMM **282** available to a subsequent SENTER instruction. In this example, software loads a copy of SINIT-AC **280** and a copy of SVMM **282** into one or more memory pages **250-262**. One processor, in the present example processor **202**, is then selected to be the ILP, which issues the SENTER instruction at time **612**. At time **614** the ILP's SENTER instruction issues the SENTER BUS MESSAGE **616**. The ILP then issues its own SENTER ACK **608** at time **618** prior to entering a wait-for-chipset-flag state at time **628**.

Each RLP, such as processor **222**, respond to the SENTER BUS MESSAGE **616** by completing the current instruction during time **620**. The RLP then issues its SENTER ACK **622** and then enters a state **634** where it waits for an SENTER CONTINUE MESSAGE.

The chipset **240** spends time **624** setting the JOINS register **272** responsive to the SENTER ACK messages observed on system bus **230**. When the JOINS register **272** contents matches the EXISTS register **270** contents, chipset **240** sets the ALL_JOINED flag **246** at time **626**.

During this time, the ILP may remain in a loop while polling the ALL_JOINED flag **246**. When the ALL_JOINED flag **246** is set, and ILP determines that the ALL_JOINED flag **246** is set at time **630**, the ILP may then issue the SENTER CONTINUE MESSAGE during time **632**. When the SENTER CONTINUE MESSAGE is broadcast on system bus **230** at time **636**, the RLPs may enter a wait-for-join state. For example, the RLP of processor **222** enters a wait-for-join state during time period **638**.

Upon issuing the SENTER CONTINUE MESSAGE, the ILP may then (in time period **640**) bring the public key of key register **244** of chipset **240** and a copy of SINIT-AC into its secure memory **208** to form a copy of the key and a copy of SINIT-AC. In another embodiment, key register **244** may contain a digest of the public key, and the actual public key may be included in, or with, the SINIT-AC. Upon authenticating the copy of SINIT-AC as described above in connection with FIG. **2**, the ILP may then actually execute the copy of SINIT-AC within secure memory **208**.

After the copy of SINIT-AC within secure memory **208** begins execution, it then (during time period **640**) validates and registers the memory-resident copy of SVMM. After the copy of SVMM is registered in the PCR **278** of security token **276**, the memory-resident copy of SVMM itself begins

execution. At this time, during ongoing time period **650**, SVMM operations are established in the ILP.

Among the first things that the ILP SVMM operation does is issue individual RLP JOIN MESSAGES on the system bus **230**. An example is a processor **222** JOIN MESSAGE **644**. This message may include a location in memory at which the RLP processor **222** may join in execution of the registered memory-resident copy of SVMM. Alternatively, the ILP SVMM operations may have registered a memory location in a predetermined location in the chipset or memory, and upon receiving the JOIN MESSAGE the RLP retrieves its starting address from this location. After receiving the processor **222** JOIN MESSAGE, and determining its starting address, during time period **646** the RLP processor **222** jumps to this location and joins execution of the registered memory-resident copy of the SVMM.

After all the RLPs have joined the registered memory-resident copy of the SVMM, secured operations are established throughout the microcomputer system **200**.

Referring now to FIG. 7, a flowchart of software and other process blocks is shown, according to one embodiment of the present invention. For the sake of clarity FIG. 7 only shows process blocks for a single representative RLP. In other embodiments there may be several responding logical processors.

The process **700** begins at block **710** when a logical processor makes a copy of the SINIT-AC and SVMM modules available for access by a subsequent SENTER instruction. In this example, in block **712** the ILP loads the SINIT-AC and SVMM code from mass storage into physical memory. In alternative embodiments, any logical processor may do so, not just the ILP. A processor becomes the ILP by executing the SENTER instruction, as identified in block **714**. In block **716**, the ILP SENTER instruction issues an SENTER BUS MESSAGE in block **716**. The ILP then, in block **718**, issues its own SENTER ACK message to the chipset. The ILP then enters a wait state, shown as decision block **720**, and waits for the chipset to set its ALL_JOINED flag.

After each RLP receives the SENTER BUS MESSAGE in block **770**, it halts execution with the end of the current instruction, and then in block **772** issues its own SENTER ACK. Each RLP then enters a wait state, shown as decision block **774**, and waits for a SENTER CONTINUE MESSAGE to arrive from the ILP.

The chipset sets the corresponding bits in the JOINS register when SENTER ACK messages are received. When the JOINS register contents equals the EXISTS register contents, the chipset sets its ALL_JOINED flag, signaling the ILP to proceed from decision block **720**.

The ILP, upon exiting decision block **720** on the YES path, then issues a SENTER CONTINUE MESSAGE in block **722**. This signals each RLP to proceed from decision block **774**. Each RLP then enters a second wait state, shown as decision block **776**, and waits for a SENTER JOIN MESSAGE.

Meanwhile the ILP, in block **724**, moves the public key of the chip set and the memory-resident copy of SINIT-AC into its own secure memory for secure execution. The ILP, in block **726**, uses the key to validate the secure-memory-resident copy of SINIT-AC, and then executes it. The execution of SINIT-AC may perform tests of the system configuration and the SVMM copy, then registers the SVMM identity, and finally begins the execution of SVMM in block **728**. As part of actions performed in block **728**, the ILP SINIT code may configure device-access page table **248** and device-access

logic **247** of memory and chipset to protect those memory pages used by the memory-resident copy of SVMM **282** from interference by non-processor devices, as shown in block **754**.

After the ILP begins execution under the control of SVMM, in block **730** the ILP sends an individual SENTER JOIN MESSAGE to each RLP. After issuing the SENTER JOIN MESSAGE, the ILP then in block **732** begins SVMM operations.

The receipt of the SENTER JOIN MESSAGE causes each RLP to leave the wait state represented by decision block **776** along the YES path, and begin SVMM operations in block **780**. The SENTER JOIN MESSAGE may contain the SVMM entry point the RLP branch to when joining SVMM operations. Alternatively, the ILP SVMM code may register the appropriate RLP entry point in a system location (for example, in the chipset), to be retrieved by the RLP upon receipt of the SENTER JOIN MESSAGE.

While various embodiments disclosed include two or more processors (either logical or physical processors), it should be understood that such multi-processor and/or multi-threaded systems are described in more detail to explain the added complexity associated with securing a system with multiple logical or physical processors. An embodiment also likely to be advantageous in less complex system may use only one processor. In some cases, the one physical processor may be multi-threading and therefore may include multiple logical processors (and accordingly have an ILP and an RLP as described). In other cases, however, a single-processor, single-threaded system may be used, and still utilize disclosed secure processing techniques. In such cases, there may be no RLP; however, the secure processing techniques still operate to reduce the likelihood that data can be stolen or manipulated in an unauthorized manner.

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A processor comprising:

hardware resources; and

a logical processor to cause execution of a secure monitor in response to a dedicated instruction, wherein the secure monitor is to control access to the hardware resources from a non-secure environment including at least one of a non-secure operating system and a non-secure application executing in a non-secure environment and from a secure environment including at least one of a secure kernel and a secure application.

2. The processor of claim 1, wherein the logical processor, in response to the dedicated instruction, is to: store a copy of the secure monitor, a digital signature of a trusted copy of the secure monitor, and a key in secure memory associated with the logical processor; authenticate the copy of the secure monitor using the digital signature and the key; and initiate execution of the secure monitor when authenticated.

3. The processor of claim 1, wherein the secure monitor is to have a higher privilege level than the non-secure operating system, the non-secure application, the secure kernel, and the secure application.